

# Chapter 4: Using the dCache to Copy Files to/from Enstore

Whenever a client application needs to talk to the dCache, it has to choose an appropriate *door* into the system. For each door, there are corresponding utilities for copying files back and forth between your machine and your `/pnfs/storage-group` area on the machine running dCache. We describe how to use the supported utilities in this chapter.

Currently (November 2003), there are four Fermilab dCache server nodes, three corresponding to Enstore installations (FNDCA, CDFDCA, and D0DCA), and CMSDCA for CMS. Each dCache server may have multiple doors, thus allowing a variety of access methods. Each door is limited to about 50 simultaneous transfers; more doors can be added as needed. The dCache supports Kerberos V5 for FTP, the dCache native dCap C-API, and GSI FTP.



The dCache server node and the ports documented in this section are subject to change. You can always find the current configuration from the web page [http://www-isd.fnal.gov/enstore/dcache\\_user\\_guide.html](http://www-isd.fnal.gov/enstore/dcache_user_guide.html).<sup>1</sup>

## 4.1 DCache-Native dCap

---

### 4.1.1 About dCap

DCap is a dCache-native access protocol. It is available in KITS at <ftp://fnkits.fnal.gov/products/dcap/>. The libdcap library provides POSIX-like **open**, **create**, **read**, **write** and **lseek** functions to the dCache storage. In addition there are some specific functions for setting debug level, getting error messages, and binding the library to a network interface. See <http://www-dcache.desy.de/manuals/libdcap.html> for usage information.

---

1. It is available from the *Fermilab Mass Storage Systems* home page (<http://hppc.fnal.gov/enstore/>); see the list of items under *Documentation* for dCache, and use the *User Access at FNAL* link.

If your dCap door uses Kerberos V5 authentication, first obtain a Kerberos principal for the FNAL.GOV realm, if you don't already have one. Install the dCap product on your computer. See [http://www-dcache.desy.de/manuals/dcap\\_setup.html](http://www-dcache.desy.de/manuals/dcap_setup.html).

The nodes and ports available for dCap are subject to change; to get a current listing, run the following command, using your storage group (sample output shown for storage group cdfen):

```
% cat '/pnfs/cdfen/.(config)(dCache)(dcache.conf)'  
  
cdfdca.fnal.gov:25125  
cdfdca.fnal.gov:25136  
...  
cdfdca2.fnal.gov:25153  
cdfdca2.fnal.gov:25154  
cdfdca3.fnal.gov:25155  
...
```

The dCap protocol requires specification of the dCache server host, port number, and domain, in addition to the inclusion of “/usr” ahead of the storage group designation in the PNFS path. Its structure is shown here:

```
dcap://<serverHost>:<port>/</pnfs>/<storage_group>/usr/<filePath>
```

There are supposed to be two slashes inbetween the port number and pnfs, e.g., ... :24124//pnfs/..., but since users frequently just put one slash, we've allowed either one or two.

## 4.1.2 The dccp Command

The command **dccp**, which provides a **cp**-like functionality on the PNFS file system, is available in the dCap product. The **dccp** command has the following syntax:

or, more simply:

```
% dccp [ options ] source_file [ destination_file ]
```

The options and command usage are described at <http://www-dcache.desy.de/manuals/dccp.html>.

A useful related command is:

```
dc_stage [-t <number of seconds>] source [ dest]
```

This prestages the request; for read requests only. It is particularly useful when you'd like to grab the file quickly from the dCache when you're ready for it. Use this with the **-t** option to set an interval of time between the download to the dCache and the download from the dCache to your local system. If **-t** is not used, the default interval is zero.



If you run a **dccp** command and it fails because the port is unavailable, try the command again with a different port number, or with a different host and port combination.

## Syntax and Examples (PNFS Not Mounted Locally)

If PNFS is not mounted locally (the general case), you'll have to supply the protocol, node, port, and pnfs directory for the remote location (the “source” on reads, and the “destination” on writes). For example, a command requesting a write to Enstore would have this structure:

```
% dccp path/to/local/file \  
dcap://<serverHost>:<port>/</pnfs>/<storage_group>/usr/<filePath>
```

Here is an example of this, requesting a write from your local /tmp directory:

```
% dccp /tmp/myfile \  
dcap://cdfdca.fnal.gov:25140//pnfs/fnal.gov/usr/cdfen/x/myfile
```

To check if a file is on disk in the dCache, run **dc\_check**:

```
% dc_check \  
dcap://fndca.fnal.gov:24124/pnfs/fnal.gov/cdf/myfile
```

If this were to be a read rather than a write, it would look like:

```
% dccp \  
dcap://cdfdca.fnal.gov:25140//pnfs/fnal.gov/usr/cdfen/x/myfile\  
/tmp/myfile
```

To pre-stage this same request with an hour interval, use **dc\_stage**:

```
% dc_stage -t 3600 \  
dcap://cdfdca.fnal.gov:25140//pnfs/fnal.gov/usr/cdfen/x/myfile\  
/tmp/myfile
```

## Syntax and Examples (PNFS Mounted Locally)

If PNFS is mounted on your local machine, you only need to specify the simple PNFS path of the remote file, e.g. (for a write):

```
% dccp path/to/local/file /pnfs/<filePath>
```

For example (using the same file as in the previous examples):

```
% dccp /tmp/myfile /pnfs/cdfen/x/myfile
```

will write the file to Enstore, and the following will read it from Enstore and put it into your local /tmp directory:

```
% dccp /pnfs/cdfen/x/myfile /tmp/myfile
```

## 4.2 Grid (GSI) FTP

---

GSI stands for Grid Security Interface. GSI FTP uses Grid Proxies for authentication and authorization and is compatible with popular Grid middleware tools such as globus-url-copy (from the Globus toolkit available at <http://www.globus.org> or from sam\_gridftp in Kits). The dCache GSI FTP currently runs on port 2811 on the following nodes (different nodes for different user groups):

- `findca.fnal.gov`, port 2811 (for general users)
- `cdafdca`, port 2811 (for CDF)
- `d0dca`, port 2811 (for D0)
- `cmsdca`, port 2811 (for CMS)

It is more convenient to run this through an interface like `srmcp` (see section 4.2.4 *Storage Resource Management (SRM)*) which allows you to perform multiple transfers in a single command. In addition, it optimizes the parameters of the transfer, and allows FTP to scale with user load (overcoming a passive gridftp protocol issue).

### 4.2.1 Obtain Grid Proxies

Globus tools require that a user be authenticated with a short-term authentication Grid proxy. This proxy can be created from (long-term) X.509 credentials issued by DOE science grid (or other Certificate Authority listed on <http://computing.fnal.gov/security/pki>) or from Kerberos credentials at Fermilab. A proxy expires after a preset duration, and then a new one must be regenerated from the user's (long-term) X.509 certificate.

X.509 Grid proxies can be issued automatically for Fermilab users authenticated to Kerberos. See

<http://computing.fnal.gov/security/pki/> for instructions.

This involves downloading a KX.509 certificate. KX.509 can be used in place of permanent, long-term certificates. It works by creating X.509 credentials (certificate and private key) using your existing Kerberos ticket. These credentials are then used to generate the Globus proxy certificate. KX.509 is described at

<http://www.ncsa.uiuc.edu/~aloftus/NMI/kx509.html>.

For non-Fermilab people, Grid proxies typically must be created from X.509 certificates. See

<http://www.doegrids.org/pages/cert-request.htm>.

## 4.2.2 GSI FTP with globus-url-copy

Install the Globus toolkit (available from a variety of locations, <http://www.globus.org> is one). Then run the **globus-url-copy** command in order to use the GSI FTP protocol to transfer files. Use the **gsiftp://** URL prefix for the PNFS (Enstore) path, and **file://** for the other URL.

E.g., to copy **from** Enstore the syntax is:

```
% globus-url-copy
  gsiftp://[<src_node>:port]/<source_url_path>
  file://[<dest_node>:port]/<dest_url_path>
```

and to copy **to** Enstore, it's:

```
% globus-url-copy file://[<src_node>:port]/<source_url_path>
  gsiftp://[<dest_node>:port]/<dest_url_path>
```

In the case of a CDF user copying from Enstore to a local disk, this would look like:

```
% globus-url-copy gsiftp://cdfdca.fnal.gov:2811/<pnfs_path>
  file:///<local_url_path>
```

A D0 user copying from a remote disk to Enstore would use a command like this:

```
% globus-url-copy file://<remotenode>:<port>/<remote_url_path>
  gsiftp://d0dca.fnal.gov:2811/<pnfs_path>
```

You can also copy from one Enstore system to another, e.g., from CDFDCA to FNDCA.

```
% globus-url-copy gsiftp://cdfdca.fnal.gov:2811/<pnfs_path>
  gsiftp://fndca.fnal.gov:2811/<pnfs_path>
```

## 4.2.3 GSI FTP with Kftpcp

GSI FTP is also available with kftpcp (see section 4.4 *Kerberized FTP via the kftpcp Command*). Install and setup kftp (from Kits <ftp://fnkits.fnal.gov/products/>). Also from kits, install and setup gsspy\_gsi (for Grid proxy) instead of gsspy\_krb. Kftpcp works the same as described in section 4.4 except that the port number is 2811 in this case.

We refer you to section 4.4 for details, but here's a quick example for a general user (using STKEN) to copy from Enstore to a local disk:

```
% kftpcp -p 2811 -m p [-v] \
  [<your_login_id>@]fndca:<pnfs_path> \ </path/to/local_file>
```

## 4.2.4 Storage Resource Management (SRM)

SRM is middleware for managing storage resources on a grid. The SRM implementation within the dCache manages the dCache/Enstore system. It provides functions for file staging and pinning<sup>1</sup>, transfer protocol negotiation and transfer url resolution.

The SRM client **srmcp** provides a convenient way to transfer multiple files from/to Enstore via dCache using a variety of protocols.

To read about SRM, go to <http://sdm.lbl.gov/>, click on **PROJECTS**, and look for *Storage Resource Management (SRM) Middleware Project*.

**Srmcp** is the implementation of SRM client as specified by the SRM spec (see <http://sdm.lbl.gov/srm/documents/joint.docs/srm.v1.0.doc>). You can use **srmcp** for the retrieval and/or storage of files to/from Enstore (or other Mass Storage Systems which implement SRM, e.g., SLAC's, CERN's). In this document we focus on file transfers to/from Fermilab's Enstore via dCache.

### Preparing to Use srmcp

Two packages are available, one with java (srmcp), the other with a C-based client (srmttools); they are both in Kits (<ftp://fnkits.fnal.gov/products/>). To use the java-based **srmcp**, you will need to install java on your system. You will also need to install either the globus toolkit or dccp, depending on which protocol you wish to use. In order to use GSI with **srmcp**, follow the instructions in the README.SECURITY file that comes with srmcp v1\_2 in Kits.

### Command Syntax

```
% srmcp [options] source(s) destination
```

Default options will be read from a configuration file but can be overridden by command line options. The options are listed and defined in the srmcp v1\_2 README file in Kits. We do not list them here.

The SRM protocol, used for the remote file specification, requires the SRM server host, port number, and domain. For the fnal.gov domain, the inclusion of “/usr” ahead of the storage group designation in the PNFS path is also required. Its structure is shown here:

```
srm://<serverHost>:<portNumber>/<root of fileSystem>  
/<storage_group>[/usr]/<filePath>
```

Some examples, the first two for the fnal.gov domain, the third for cern.ch:

---

1. Pinning refers to making a file undeletable in the cache for the period of time called the “lifetime of the job”.

- `srms://cdfdca.fnal.gov:25129//pnfs/fnal.gov/usr/cdfen/filesets/<filePath>`
- `srms://fnisd1.fnal.gov:24129//pnfs/fnal.gov/usr/cdfen/<filePath>`
- `srms://wacdr002d.cern.ch:9000/castor/cern.ch/user/<filePath>`

## Examples

These examples are taken from the srmcp v1\_2 README file in Kits (with unnecessary options removed).

The following command will retrieve two files `/mypath/myfile1.ext` and `/mypath/myfile2.ext` from Enstore via dCache (for a CDF user) and store them in the user's local directory `/home/me/targetdir`. Notice that srmcp requires that the PNFS path include `/pnfs/fnal.gov/usr/` ahead of the storage group designation.

```
% srmcp \
  srm://cdfdca.fnal.gov:25129//pnfs/fnal.gov/usr/cdf/myfile1.ext \
  srm://cdfdca.fnal.gov:25129//pnfs/fnal.gov/usr/cdf/myfile2.ext \
  file://localhost/home/me/targetdir
```

The following will copy the same files from one Enstore installation (CDFEN) to another (STKEN):

```
% srmcp \
  srm://cdfdca.fnal.gov:25129//pnfs/fnal.gov/usr/cdf/myfile1.ext \
  srm://cdfdca.fnal.gov:25129//pnfs/fnal.gov/usr/cdf/myfile2.ext \
  srm://fndca.fnal.gov:24128/targetdir
```

The following will get the file using dccp client, overriding the default (dccp would have to be already installed on you machine)<sup>1</sup>:

```
% srmcp \
  -protocols=dcap \
  srm://fndca.fnal.gov:24128//pnfs/fnal.gov/usr/targetdir/myfile1.ext
file:///tmp/myfile1.ext\
```

## 4.3 Simple Kerberized FTP

---

The dCache door for Kerberized ftp service enforces Kerberos authentication (see *Strong Authentication at Fermilab Documentation* at <http://computing.fnal.gov/docs/strongauth/>). It currently runs on the following nodes and corresponding ports:

- `fndca.fnal.gov`, port 24127 (for STKEN)

---

1. The four slashes in the last line refer to: `file://`; host, which comes next, is “”; path is `/tmp/...`

- cdfdca, port 25127 (for CDFEN)
- d0endca, port 24127 (for D0EN)

(The port number is installation-specific.) Any Kerberized ftp client can be used on the client machine. You must specify the host port in your ftp command.

Notes:

- File read and write functionality is supported when the user (a) is authorized by the experiment to access the data stores, and (b) has obtained Kerberos credentials.
- Portal Mode (CRYPTOCARD) access is not supported since it is not compatible with automated transfers or future GRID development.

### 4.3.1 Prepare to use Kerberized FTP

In order to establish the kftp service on dCache, you must first:

- have a valid Fermilab UNIX account (UID and GID)
- have a Kerberos principal for FNAL.GOV (if Kerberized access is required)
- ask your experiment's Enstore liaison to register you for the service; you'll need to provide the following information to the liaison:
  - username
  - UID and GID (run the command `id` at the UNIX prompt to find their values)
  - storage group
  - root path under `/pnfs/<storage_group>/...`
  - if applying for Kerberized door, provide Kerberos principal(s)
  - if applying for weak door, request a password by emailing `dcache-admin@fnal.gov`. This is for groups, not individuals.
- install the **kftp** product from KITS (optional; useful for running scripts to transfer files). To do so, run:
 

```
$ setup upd
$ upd install -G "-c" kftp
```

### 4.3.2 Sample Kerberized FTP session

User is authenticated to Kerberos and authorized for the Kerberized dCache door (currently at `fndca.fnal.gov`, port 24127):

```
% ftp fndca.fnal.gov 24127
Connected to stkendca3a.fnal.gov.
```



```

220 FTPDoorIM+GSS ready
334 ADAT must follow
GSSAPI accepted as authentication type
GSSAPI authentication succeeded
Name (fndca:acheavey):
200 User acheavey logged in
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> cd acheavey/test3
250 CWD command successful. New CWD is </acheavey/test3>
ftp> ls
200 PORT command successful
150 Opening ASCII data connection for file list
dupl2
duplexps
226 ASCII transfer complete
ftp> get duplexps
local: duplexps remote: duplexps
200 PORT command successful
150 Opening BINARY data connection for /pnfs/fs/usr/test/acheavey/test3/duplexps
226 Closing data connection, transfer successful
42 bytes received in 0.033 seconds (1.2 Kbytes/s)
ftp>

```

## 4.4 Kerberized FTP via the **kftpcp** Command

---

In order to access data from a batch job or a background process, you should either use ftp client libraries (available from many sources), or the **kftp** package. This package includes a Kerberized client library and a GSI client library; you can use either. A regular ftp client (Kerberized or not) is an interactive program which is hard to use in batch mode.

See section 4.3.1 *Prepare to use Kerberized FTP* for installation information. To use the product in a UPS environment as a Kerberized FTP client, first run:

```
% setup gsspy_krb; setup kftp
```

Then run the **kftpcp** command to copy one or more files. This command can be used from the shell or in a script.

### 4.4.1 Syntax and Options

```
% kftpcp [<options>] <source_file> <destination_file>
```

The available options include:

<b>-p</b> <port>	ftp server port number
<b>-m</b> <a   p>	ftp server mode; active (default), or passive
<b>-v</b>	verbose mode

Notes:

- If your login id is the same on fndca and your local system, and if they match your Kerberos principal, you can leave off `<your_fndca_login_id>@` in front of `fndca:` in the command.
- Depending on how your access is configured, typically you only need to specify the path to the remote file starting from the directory under your `/pnfs/<storage_group>/` area. E.g., to specify the remote file `/pnfs/my_storage_group/path/to/file` on the command line, enter only `/path/to/file`, including the initial slash. You can use the full specification (starting with `/pnfs/<domain>/usr/<storageGroup>`)

## 4.4.2 Download a File

To download a stored data file from Enstore via the dCache, using fndca as a sample server host, run:

```
% kftpcp -p 24127 -m p [-v] \
  [<your_fndca_login_id>@]fndca:</path/to/remote_file> \
  </path/to/local_file>
```

## 4.4.3 Upload a File

To upload a new data file, again using fndca, run:

```
% kftpcp -p 24127 -m p [-v] \
  </path/to/local_file> \
  [<your_fndca_login_id>@]fndca:</path/to/remote_file>
```

## 4.4.4 Examples

To read (download) the stored file `/pnfs/storage_group/mydir/myfile` into a local file of the same name, run:

```
% setup kftp
% kftpcp -p 24127 -m p -v myloginid@fndca:/mydir/myfile \
  /path/to/myfile
```

Transferred 42 bytes

Or, if your usernames and principal all match, you could shorten it to:

```
% kftpcp -p 24127 -m p -v fndca:/mydir/myfile /path/to/myfile
```

## 4.5 Weakly-Authenticated FTP Service (Read-only)

---

The dCache weakly-authenticated ftp service currently runs on node the following nodes and corresponding ports:

- fndca.fnal.gov, port 24126 (for STKEN).
- cdfdca, port 25126 (for CDFEN)
- d0endca, port 24126 (for D0EN)

This is read-only, and is not necessarily allowed by all experiments. This ftp service can be accessed by ordinary ftp client software. You must specify the host port in your ftp command, as shown below. The Enstore admin will have sent you an email to confirm your registration for this service, and included a password for it.<sup>1</sup> This is a weak password. Log in with your username and password.

### Sample weakly-authenticated read-only ftp session

Here we explicitly use a weakly-authenticated ftp client, `/usr/bin/ftp`, and make the connection to fndca port 24126. In the session, we first successfully retrieve a file called `myfile`, and secondly attempt to write a file `trace.txt` and (correctly) fail.

```
% /usr/bin/ftp fndca.fnal.gov 24126

Connected to stkendca3a.fnal.gov.
220 FTPDoorIM+PWD ready (read-only server)
Name (fndca:acheavey):
331 Password required for acheavey.
Password: (password entered here)
230 User acheavey logged in
ftp> cd acheavey/test3
250 CWD command successful. New CWD is </acheavey/test3>
ftp> ls
200 PORT command successful
150 Opening ASCII data connection for file list
myfile
myfile2
myfile3
226 ASCII transfer complete
10 bytes received in 0.018 seconds (0.55 Kbytes/s)
ftp> get myfile
200 PORT command successful
150 Opening BINARY data connection for
    /pnfs/fs/usr/test/acheavey/test3/myfile
226 Closing data connection, transfer successful
local: myfile remote: myfile
42 bytes received in 0.05 seconds (0.82 Kbytes/s)
ftp> put trace.txt
200 PORT command successful
500 Command disabled
```

---

1. If you need to change this password, send email to [dcache-admin@fnal.gov](mailto:dcache-admin@fnal.gov).

```
ftp> bye
```